

Beyond $O^*(2^n)$ in domination-type problems

Marek Cygan* Marcin Pilipczuk* Jakub Onufry Wojtaszczyk†

September 22, 2009

Abstract

In this paper we provide algorithms faster than $O^*(2^n)$ for several NP-complete domination-type problems. More precisely, we provide:

- an algorithm for CAPACITATED DOMINATING SET that solves it in $O(1.89^n)$,
- a branch-and-reduce algorithm solving LARGEST IRREDUNDANT SET in $O(1.9657^n)$ time,
- and a simple iterative-DFS algorithm for SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET that solves it in $O(1.999956^n)$ time.

We also provide an exponential approximation scheme for CAPACITATED DOMINATING SET. All algorithms require polynomial space. Despite the fact that the discussed problems are quite similar to the DOMINATING SET problem, we are not aware of any published algorithms solving these problems faster than the obvious $O^*(2^n)$ solution prior to this paper.

1 Introduction

Notation and definitions. Let $G = (V, E)$ be an undirected graph. Given $F \subset E$ we write $V(F)$ to denote the set of all endpoints of the edges in F . Given $W \subset V$ by $G[W]$ we denote the subgraph induced by W .

We say a vertex $v \in V$ dominates $u \in V$ if $u = v$ or $uv \in E$, i.e. a vertex dominates itself and all its neighbours. By $\bar{N}(v) = \{v\} \cup \{u : uv \in E\}$ we denote the set of vertices dominated by v . We extend this notation to any subset $W \subset V$ by putting $\bar{N}(W) = \bigcup_{v \in W} \bar{N}(v)$, we say that a set W dominates a vertex u if $u \in \bar{N}(W)$. The set $\bar{N}(W)$ is called the *closed neighbourhood* of W .

The DOMINATING SET problem asks for the smallest set that dominates the whole V . In the CAPACITATED DOMINATING SET problem each vertex v is additionally equipped with a number $c(v)$, which is the number of other vertices this vertex can dominate. Formally, we say that a set $S \subset V$ is a capacitated dominating set if there exists $f_S : V \setminus S \rightarrow S$ such that $f_S(v)$

*Dept. of Mathematics, Computer Science and Mechanics, University of Warsaw, Poland

†Institute of Mathematics, Polish Academy of Sciences

is a neighbour of v for each $v \in V \setminus S$ and $|f_S^{-1}(w)| \leq c(w)$ for each $w \in S$. The function f_S is called a *dominating function* for the set S . The CAPACITATED DOMINATING SET problem asks for the smallest possible size of a capacitated dominating set. Note that for a given set S checking whether it is a capacitated dominating set is a polynomial-time problem which can be solved using max-flow or maximum matching techniques.

We say a set $S \subset V$ is *irredundant* if for any $v \in S$ there exists a vertex $u \in V$ such that v dominates u and $S \setminus \{v\}$ does not dominate u . We call any such vertex u a *unique* vertex for v . An irredundant set is called *inclusion-maximal* if it is not a proper subset of any other irredundant set. Note that an inclusion-maximal irredundant set does not necessarily have to dominate the whole vertex set of G , for examples see [9].

The irredundance numbers $ir(G)$ and $IR(G)$ are defined as the smallest and largest possible cardinality of an inclusion-maximal irredundant set. The problems SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET and LARGEST IRREDUNDANT SET ask for $ir(G)$ and $IR(G)$ respectively.

Our results. In Section 2 we provide an algorithm which solves the CAPACITATED DOMINATING SET problem in $O(1.89^n)$. The algorithm constructs $O^*(\binom{n}{n/3}) = O(1.89^n)$ reductions of the input graph into a SIMPLIFIED CAPACITATED DOMINATING SET problem instance (defined in Section 2.1), each solvable in polynomial time. Section 2.3 tackles the exponential approximation of CAPACITATED DOMINATING SET. More precisely, we provide an approximation algorithm that for a given $c \in (0, \frac{1}{3})$, in time $O^*(\binom{n}{cn}) = O^*((1/(c^c(1-c)^{1-c}))^n)$, computes a $(\frac{1}{4c} + c)$ -approximation in the case of $c < \frac{1}{4}$ or a $(2 - 3c)$ -approximation in the case of $\frac{1}{4} \leq c < \frac{1}{3}$.

In Section 3 we provide a branch-and-reduce algorithm that solves LARGEST IRREDUNDANT SET in $O(1.9657^n)$. In Section 4 we consider simple iterative depth-first search algorithm that solves SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET and prove that it breaks slightly the $O^*(2^n)$ barrier: we show that it works in $O(1.999956^n)$ time.

All algorithms in this paper require polynomial space.

Motivation and previous work. The field of exact exponential-time algorithms for NP-hard problems has attracted a lot of attention in the recent years (see Woeginger's survey [26]). Many difficult problems can be solved much faster than by the obvious brute-force algorithm; examples are INDEPENDENT SET [18], DOMINATING SET [17, 25], CHROMATIC NUMBER [2] or BANDWIDTH [11, 15]. A few powerful techniques have been developed, including Measure & Conquer [17] and inclusion/exclusion principle applications [2, 3, 22]. However, there is still a bunch of problems for which no faster solution than the obvious one is known. These include SUBGRAPH ISOMORPHISM and CHROMATIC INDEX which are mentioned as open problems in [1, 19].

Superpolynomial approximation was recently considered as a way of coping with hardness of approximation of different NP-hard problems. Results in this field include subexponential approximation algorithm for BANDWIDTH on trees [12, 16] and exponential approximation schemes for CHROMATIC NUMBER [7, 10] or BANDWIDTH on arbitrary graphs [12, 20].

Finding an algorithm faster than $O^*(2^n)$ for DOMINATING SET was an open problem until 2004. Currently the fastest algorithm by van Rooij et al. runs in $O(1.5048^n)$ [25]. However, many other graph-dominating problems still lack faster solutions. The problems of solving

CAPACITATED DOMINATING SET, SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET and LARGEST IRREDUNDANT SET faster than obvious $O^*(2^n)$ algorithms were posted by van Rooij in 2008 [19], CAPACITATED DOMINATING SET was also posted as an open problem at IWPEC 2008. In this paper we present algorithms providing answers to those problems.

Note that at first glance breaking $O^*(2^n)$ barrier for the CAPACITATED DOMINATING SET problem seems a hard task, since even the brute-force $O^*(2^n)$ algorithm involves matching or max-flow techniques. Moreover, from the parameterized point of view, Dom et al. [13] showed that this problem is $W[1]$ -hard when parameterized by both treewidth and solution size, and Bodlaender et al. [6] showed that even the planar version parameterized by the solution size is also $W[1]$ -hard.

Our algorithm for CAPACITATED DOMINATING SET is somewhat similar to one of the first algorithms to break $O^*(2^n)$ for the classical DOMINATING SET problem, namely the algorithm of Randerath and Schiermeyer [24]. Their algorithm also involves matching arguments and our algorithm, applied to DOMINATING SET, can be viewed as a simplification of their algorithm. However we do not know whether their algorithm could be used to solve the CAPACITATED DOMINATING SET problem.

Downey et al. [14] showed that the problem of finding an irredundant set of size k is $W[1]$ -hard when parameterized by k and FPT when parameterized by $n - k$. However, Raman and Saurabh proved that this problem become FPT in graphs without small cycles [23]. It is worth noticing that irredundant sets attract some attention from the graph-theoretical point of view [9].

There is a large number of problems for which the fastest current known algorithm requires $O^*(2^n)$ time, including CHROMATIC NUMBER and TSP. Koivisto [19] conjectured that there might exist some negative results, i.e. reductions in the following manner: if one problem can be solved in $O^*(c^n)$ time for some $c < 2$, then another problem can be solved in $O^*(d^n)$ time for some $d < 2$. This paper shows that the problems considered by us are somewhat easier and admit solutions faster than $O^*(2^n)$. We would like to note that our techniques in Section 4 seem a bit similar to those that lead to $O^*((2 - \varepsilon)^n)$ -time algorithms for DOMINATING NUMBER [5] and TSP [4] in graphs of bounded degree. We think it is interesting and somewhat surprising that such techniques can be used in graphs without any degree assumption.

Recently it came to our attention that Brankovic et al. [8] were independently working on the irredundance problems but we are not aware of the details of their results.

Acknowledgements We would like to sincerely thank George R. R. Martin for his Song of Ice and Fire novels and O. S. Card for his Ender’s series — they have given us a number of enjoyable and entertaining breaks while we were working on this paper.

2 CAPACITATED DOMINATING SET

2.1 SIMPLIFIED CAPACITATED DOMINATING SET

In this section we introduce a simplification of the CAPACITATED DOMINATING SET problem, namely the SIMPLIFIED CAPACITATED DOMINATING SET problem, which can be solved in polynomial time.

The input of **SIMPLIFIED CAPACITATED DOMINATING SET** is an undirected graph $G = (V, E)$, a set $U \subset V$ and a capacity function $c : V \rightarrow \mathbb{N}$. We ask for the smallest capacitated dominating set $S \subset G$ containing U such that there exists a dominating function f_S satisfying

$$|f_S^{-1}(v)| \leq 1 \quad \text{for each } v \in V \setminus U. \quad (2.1.1)$$

Let $G = (V, E)$ with $U \subset V$ and a capacity function $c : V \rightarrow \mathbb{N}$ be a **SIMPLIFIED CAPACITATED DOMINATING SET** instance. Consider the following graph $G' = (V', E')$:

- for any $v \in V \setminus U$ we have $v \in V'$;
- for any $v \in U$ we have $c(v)$ copies $v_1, v_2, \dots, v_{c(v)}$ of v in V' ;
- for any $v \in V \setminus U$ and $u \in U$ the edge $uv_i \in E'$ for all i iff $uv \in E$;
- for $v, w \in V \setminus U$ we have $vw \in E'$ iff $vw \in E$ and $c(v) + c(w) > 0$;
- there are no edges of the form $v_i w_j$ or $v_i v_j$ for $v, w \in U$.




Figure 1: From the simple capacitated dominating set $\{A, B, C, D, L\}$ to a matching.

We show a correspondence between feasible solutions of **SIMPLIFIED CAPACITATED DOMINATING SET** in G and matchings in G' .

Lemma 2.1. *Let S be a capacitated dominating set in G with dominating function f_S satisfying condition 2.1.1. Then one may construct in polynomial time a matching $\phi(S, f_S)$ in G' satisfying $|V| - |\phi(S, f_S)| = |S|$.*

Proof. Let us define the matching $\phi(S, f_S)$ as follows:

- for each $v \notin S$ such that $f_S(v) \notin U$ add $vf_S(v)$ to $\phi(S, f_S)$;
- for each $v \notin S$ such that $u = f_S(v) \in U$ add vu_i to $\phi(S, f_S)$, where u_i is a copy of u in G' and different copies u_i are chosen for different vertices v with $f_S(v) = u$ (note that $|f_S^{-1}(u)| \leq c(u)$, so there are enough vertices u_i).

Note that every vertex $v \in V \setminus S$ is an endvertex of an edge in the matching $\phi(S, f_S)$. The second endvertex is $f_S(v)$ (in the case $f_S(v) \notin U$) or a copy of $f_S(v)$ in G' (in the case $f_S(v) \in U$). Moreover, every edge in $\phi(S, f_S)$ has an endpoint in $V \setminus S$. Therefore $|\phi(S, f_S)| + |S| = |V|$. \square

Lemma 2.2. *Let M be a matching in G' . Then one may construct in polynomial time a feasible solution $\psi(M)$ to the SIMPLIFIED CAPACITATED DOMINATING SET problem with dominating function $f_{\psi(M)}$ satisfying $|\psi(M)| = |V| - |M|$.*

Proof. Consider the following capacitated dominating set $\psi(M)$ with dominating function $f_{\psi(M)}$:

- $U \subset \psi(M)$;
- for $u \in U$ and for each i such that $u_i v \in M$, we take $f_{\psi(M)}(v) = u$;
- for any edge $vw \in M$, where $v, w \notin U$ one of the endpoints (say v) has to satisfy $c(v) > 0$, we add v to $\psi(M)$ and set $f_{\psi(M)}(w) = v$;
- for any $v \notin U$ which is not an endpoint of any edge in M we add v to $\psi(M)$.

It is easy to verify that the above procedure does indeed give a feasible solution to SIMPLIFIED CAPACITATED DOMINATING SET. We have $|\psi(M)| = |V| - |M|$ since for each edge in M , exactly one of its endpoints does not belong to $\psi(M)$. \square

We conclude this section with the following theorem.

Theorem 2.3. *The SIMPLIFIED CAPACITATED DOMINATING SET problem can be solved in polynomial time.*

Proof. By Lemmas 2.1 and 2.2, to find the solution of the SIMPLIFIED CAPACITATED DOMINATING SET problem it is enough to find any maximum matching in G' , which can be done in polynomial time (see e.g. [21]). \square

2.2 From SIMPLIFIED CAPACITATED DOMINATING SET to CAPACITATED DOMINATING SET

Let us start with the following simple observation. Let S be any capacitated dominating set and let f_S be a dominating function for S . Let

$$U_S = \{v \in S : |f_S^{-1}(v)| \geq 2\}.$$

We have

$$\sum_{v \in S} 1 + |f_S^{-1}(v)| = |S| + \sum_{v \in S} |f_S^{-1}(v)| = n,$$

thus in particular $|U_S| \leq n/3$. Moreover, S with the function f_S is a feasible solution for the SIMPLIFIED CAPACITATED DOMINATING SET instance with the graph G and the set U_S . Therefore the following algorithm solves CAPACITATED DOMINATING SET:

1. For each $U \subset V$ satisfying $|U| \leq n/3$ solve the SIMPLIFIED CAPACITATED DOMINATING SET instance with graph G and subset U .
2. Return the smallest capacitated dominating set from the constructed SIMPLIFIED CAPACITATED DOMINATING SET instances.

The SIMPLIFIED CAPACITATED DOMINATING SET problem can be solved in polynomial time and there are

$$\sum_{k=0}^{\lceil n/3 \rceil} \binom{n}{k} = O^* \left(\binom{n}{\lceil n/3 \rceil} \right) = O(1.89^n)$$

possible sets U (i.e. sets of cardinality at most $n/3$), thus the whole algorithm works in $O(1.89^n)$ time.

2.3 Approximating CAPACITATED DOMINATING SET

It is known that DOMINATING SET is as hard to approximate as SET COVER and since CAPACITATED DOMINATING SET problem is not easier there probably does not exist a polynomial time algorithm solving CAPACITATED DOMINATING SET with a constant approximation ratio. If we do not have enough time to obtain an exact solution for the CAPACITATED DOMINATING SET problem we can use the following constant approximation scheme. Instead of investigating all subsets $U \subset V$ satisfying $|U| \leq n/3$ we can check only smaller sets, namely $|U| \leq cn$ for some constant $c \in (0, \frac{1}{3})$. Thus the approximation algorithm has the following form:

1. For each $U \subset V$ satisfying $|U| \leq cn$ solve the SIMPLIFIED CAPACITATED DOMINATING SET instance with graph G and subset U .
2. Return the smallest capacitated dominating set from the constructed SIMPLIFIED CAPACITATED DOMINATING SET instances.

Theorem 2.4. *For any fixed constant $c \in (0, \frac{1}{3})$ the described algorithm runs in $O^* \left(\binom{n}{cn} \right) = O^* \left((1/(c^c(1-c)^{1-c}))^n \right)$ time and polynomial space. For $c \leq 1/4$ the approximation ratio is at most $(\frac{1}{4c} + c)$ and for $c \geq 1/4$ the approximation ratio is at most $2 - 3c$.*

Proof. For each subset $U \subset V$ the algorithm uses polynomial time only, thus the time bound follows directly from the Stirling formula which can be used to bound the number of subsets $\binom{n}{cn}$.

To calculate the approximation ratio let us consider some optimal solution $OPT \subset V$ together with a function $f_{OPT} : V \setminus OPT \rightarrow OPT$. By OPT_0 , OPT_1 and OPT_2 let us denote subsets of OPT containing vertices which dominate exactly zero, exactly one and at least two vertices from $V \setminus OPT$, according to f_{OPT} , respectively. By m, m_0, m_1, m_2 we denote the cardinalities of sets OPT, OPT_0, OPT_1 and OPT_2 respectively.

We may assume that $m_2 > cn$ since otherwise our algorithm finds the optimal solution. By s let us denote the average number of vertices from $V \setminus OPT$ which a vertex from OPT_2 dominates, i.e. $s = (n - m - m_1)/m_2$. Since our algorithm checks all subsets $U \subset V$ satisfying $|U| \leq cn$ it obviously considers the subset $U_0 \subset OPT_2$ containing vertices which dominate the largest number of vertices from $V \setminus OPT$. For this particular subset U_0 note that there exists a feasible solution to SIMPLIFIED CAPACITATED DOMINATING SET of size $m + (m_2 - cn)(s - 1)$: we take OPT and for each vertex in $v \in OPT_2 \setminus U_0$ we take all but one vertices from $f_{OPT}^{-1}(v)$.

Thus the approximation ratio can be bounded by $\alpha = 1 + (m_2 - cn)(s - 1)/m$. Note that if we keep m_2 fixed and increase m_0 and m_1 , the approximation ratio decreases — we increase m

and decrease s — therefore w.l.o.g. we may assume $m_0 = m_1 = 0$. Denoting $x = n/m_2$, we obtain $\alpha \leq 1 + (1 - cx)(x - 2)$ for $3 \leq x \leq \frac{1}{c}$. The bound for α is a concave function of x with maximum at $x_0 = \frac{1}{2c} + 1$. This gives $\alpha \leq \frac{1}{4c} + c$ for $c \leq \frac{1}{4}$ and $\alpha \leq 2 - 3c$ for $\frac{1}{4} \leq c \leq \frac{1}{3}$. \square

Corollary 2.5. *For $c = 1/6$ we obtain a $5/3$ -approximation algorithm which runs in $O(1.57^n)$ time and polynomial space.*

This result should be compared to the trivial approximation scheme that works in $O^*\left(\binom{n}{cn}\right)$ time too: iterate over all subsets of V that have size at most cn or at least $(1 - c)n$ and return the smallest feasible solution found. However, this algorithm has an approximation factor of $\frac{1}{c} - 1$, which is between $2\times$ and $4\times$ worse than our ratio.

3 LARGEST IRREDUNDANT SET

In this section we show a branching algorithm to find the largest (in the cardinality sense) irredundant set. Notice that the largest irredundant set is necessarily inclusion-maximal.

3.1 Bipartite independent set reduction

Let us consider a bipartite graph $H = (W, F)$, with $W = V \cup V'$ (where V' is a disjoint copy of V) and edges $\{u, v'\} \in F$ iff $\{u, v\} \in E$ or $u = v$. We call a set of edges $M \subset F$ in this graph *independent* if:

- no two edges in M share an endpoint
- the set $W(M)$ is an independent set in the graph $(W, F \setminus M)$, i.e. no edge connects endpoints of different edges from M .

Now we show a correspondence between largest irredundant sets in G and largest independent sets of edges in H :

Lemma 3.1. *If M is an independent set of edges in H , then $S := W(M) \cap V$ is an irredundant set in G . Conversely, if S is an irredundant set in G , then there exists an independent set of edges $M \subset F$ such that $S = W(M) \cap V$. In both cases, the independent set of edges M and the irredundant set S are of the same cardinality.*

Proof. Let T be an independent set of edges in H . If $\{u, v'\} \in T$, then v' is not a neighbour of any $w \in T \cap V$ other than u , thus v is a unique vertex dominated by u in G . Conversely, if we have an irredundant set S in G then letting $u(v)$ to be any unique vertex dominated by v for $v \in S$ we obtain an independent set of edges $M := \{\{v, u'(v)\} : v \in S\}$ in H . \square

Thus, to find the largest irredundant set in G it suffices to look for the largest independent set of edges in H .

A remark to make here is that the correspondence between independent edge-sets of H and irredundant sets of G does not preserve inclusion-maximality — e.g. in the graph in Figure 2 there are four inclusion-maximal irredundant sets ($\{v_1, v_3\}$, $\{v_1, v_4\}$, $\{v_2, v_3\}$ and $\{v_2, v_4\}$),




Figure 2: From the irredundant set $\{v_1, v_3\}$ to the independent set of edges $\{v_1v'_1, v_3v'_4\}$.

while in the corresponding bipartite graph the edge $v_2v'_3$ forms a singleton inclusion-maximal independent edge-set. This does no harm in the search for $IR(G)$ (as the cardinality of the set is preserved), but prevents us from using this approach to tackle the problem of $ir(G)$.

3.2 Branch and reduce analysis

Let $\alpha = 1.40202$. We provide an algorithm for finding the largest independent set of edges in a bipartite graph H on n vertices with complexity $O(\alpha^n)$. As we reduced the problem of finding the largest irredundant set in a graph with n vertices to the problem of finding the largest independent edge set in a bipartite graph with $2n$ vertices, our algorithm allows us to find the largest irredundant set in a graph on n vertices in $O(\alpha^{2n}) = O(1.9657^n)$ time.

The algorithm is a typical branch-and-reduce algorithm. We provide a set of reduction rules; in a single step the algorithm finds the first applicable rule and applies it. By $T(n)$ we denote the bound on the number of steps that the algorithm requires for a graph with n vertices. For each rule, we provide a recursion inequality for $T(n)$. Since the function $T(n) = C \cdot \alpha^n$ satisfies all inequalities for any $C > 0$, and every step takes polynomial time, the algorithm works in desired time. We do simple reductions without any memoization, so the algorithm uses polynomial space. If we can spare more memory, we can have a standard trade-off between memory and computation time: we may for instance precompute answers for subgraphs of our bipartite graph having up to $n/11$ vertices in $O^*(\binom{n}{n/11}) = O(1.36^n)$ time and space, and then stop branching when we reach the level $n/11$, this reduces the time complexity to $O(\max(1.36^n, \alpha^{10n/11})) = O(1.36^n)$ for finding the independent set of edges, and thus to $O(1.36^{2n}) = O(1.85^n)$ for finding the largest irredundant set.

Let $H = (W, F)$ be a bipartite graph on n vertices. We use the word *choose* (an edge) to mean “consider as a part of the independent edge set being built in the considered subcase” and the word *drop* (a vertex) to mean “consider that this vertex is not the endpoint of any edge in the independent edge set being built in the considered subset, so we can consider the graph without this vertex in this subcase”.

1. There is an isolated vertex $v \in W$. Obviously it cannot be an endpoint of an edge, so we drop it and solve the problem for $(W \setminus \{v\}, F)$. The recursion inequality for $T(n)$ is $T(n) \geq T(n - 1)$, which holds trivially.
2. There is a vertex v of degree 1 in H , and its only neighbour is also of degree 1. These two vertices form an edge which is a part of any largest independent edge set. Thus we choose it and solve the remaining graph on $n - 2$ vertices, obtaining the trivial inequality $T(n) \geq T(n - 2)$.

3. There is a vertex v of degree 1 in H . Let u be its sole neighbour, $\deg u \geq 2$. Consider any independent edge set M which contains some edge ux for $x \neq v$. Then $(M \setminus \{ux\}) \cup \{uv\}$ is also an independent set of the same cardinality. Thus we can branch out into two cases — either we drop u (and consequently drop v , as it becomes isolated), or we choose uv and drop all the neighbours of v (at least one of them). Here we obtain inequality $T(n) \geq T(n-2) + T(n-3)$.
4. There is a vertex v of degree $k := \deg v \geq 8$. We branch out, and either drop this vertex and obtain a graph with $n-1$ vertices, or choose one of the edges incident with v . If we choose an edge, we remove all the neighbours of v (k of them) and all the neighbours of the other endpoint (at least 2 of them — we already excluded vertices of degrees 0 and 1, and they do not coincide with the neighbours of v due to bipartity) from the graph. Thus we obtain inequality

$$T(n) \geq T(n-1) + k \cdot T(n-k-2) \quad \text{for } k \geq 8.$$

Now we check if $T(n) = C\alpha^n$ satisfies this inequality for all $k \geq 8$. However, note that $\alpha > \frac{k+1}{k}$ for $k \geq 8$, so we only need to check the case $k = 8$, which holds by straightforward calculations.

5. There are two adjacent vertices u and v of degree 2 each. Let u_1 be the other neighbour of u and v_1 be the other neighbour of v ($u_1 \neq v_1$ due to bipartity). We claim that there exists a largest independent edge set which contains one of the edges uu_1 , uv or vv_1 . Indeed, consider any independent edge set M . Then:
 - if neither u_1 nor v_1 is an endpoint of an edge in M , we may add uv to M , preserving independence and increasing cardinality;
 - if, say, u_1 is the endpoint of some edge u_1x , but none of the three aforementioned edges belong to M , we may remove u_1x from M and add u_1u instead, thus preserving cardinality and keeping independence.

Thus we consider three cases (each case consisting of choosing one of the three edges). In each case we remove the neighbours of two adjacent vertices, each of degree at least two, thus we remove at least 4 vertices in each case. We get the inequality $T(n) \geq 3T(n-4)$.

6. There is a vertex v of degree 2. Let u and w be its two neighbours. Let d_u and d_w be degrees of u and w respectively and let k be the number of common neighbours of u and w , excluding v . We have $3 \leq d_u, d_w < 8$ and $0 \leq k < d_u, d_w$. We branch into four cases. In the first, the edge uv is chosen, we drop the neighbours of u and v and obtain graph on $n - d_u - 2$ vertices. In the second, we choose the edge vw , and similarly obtain $n - d_w - 2$ vertices. In the third, we drop the vertices u, v and w and solve the problem for the graph on $n - 3$ vertices. In the fourth we drop v and choose one of the edges coincident with w and one of the edges coincident with u .

The case we have not considered is when there is an independent edge set M , where exactly one of the points w, u , say w , is coincident with an edge wx from M . In this case,

however, we may remove wx from M , adding wv instead, preserving both cardinality and independence. Thus this case need not be considered.

Now let us analyze the fourth branching option. If we choose edges uu' and ww' to the independent set of edges, neither u' nor w' is a common neighbour of u and w . Therefore we can choose u' and w' in $(d_u - k - 1)(d_w - k - 1)$ ways. Let us check how many vertices are removed in each choice. From one side of the bipartite graph we remove at least v and the $d_u - 1 + d_w - 1 - k$ vertices adjacent to either u or w . From the other we remove u , w and at least one neighbour of u' (w' and u' may both have degree 2 and common neighbours, thus we cannot assume removing additionally a neighbour of w'). Thus we remove at least $d_u + d_w + 2 - k$ vertices in total and obtain the inequality

$$T(n) \geq T(n - d_u - 2) + T(n - d_w - 2) + T(n - 3) \\ + (d_u - k - 1)(d_w - k - 1)T(n - d_u - d_w - 2 + k).$$

A direct check of all the possible values of d_u , d_w and k proves the inequality for $T(n) = C\alpha^n$.

7. There are two adjacent vertices u and v , both of degree 3. This case is the most complicated and is considered separately in Section 3.3.
8. What remains is a case where there is a vertex v of degree $3 \leq i < 8$ and all its neighbours have degree at least 4. We branch, either dropping v or choosing one of the coincident edges. If we choose edge vv' , we remove at least $i + 4$ vertices, thus obtaining the inequality $T(n) \geq T(n - 1) + iT(n - i - 4)$. By a direct check the inequality holds for $T(n) = C\alpha^n$ for all valid values of i .




Figure 3: Notation in Cases 6 and 7 (i.e. Lemma 3.2).

3.3 Two adjacent vertices of degree 3

Recall the single case we have left to consider is when we have two adjacent vertices of degree 3, and all other vertices have degrees between 3 and 7. Let the two adjacent vertices be u and v , let u_1 and u_2 be the two other vertices adjacent to u and let v_1 and v_2 be the two other

vertices adjacent to v . Again, we want to take advantage of the fact that searching for the largest independent edge set we do not have to consider some of the combinations.

Lemma 3.2. *There exists a largest independent edge set M in H for which either*

1. *one of the edges $uv, uu_1, uu_2, vv_1, vv_2$ appears in M , or*
2. *both u_1 and u_2 are endpoints of edges in M , or*
3. *both v_1 and v_2 are endpoints of edges in M and u_1 and u_2 are not endpoints of edges in M .*

Proof. Take any largest independent edge set M' . If M' satisfies the conditions above, there is nothing to prove.

If exactly one of u_1, u_2 is an endpoint of an edge in M' (say u_1 is an endpoint of $u_1u'_1$), we may remove $u_1u'_1$ from M' and add uu_1 instead. We do not lose independence (as neither v nor u_2 was an endpoint of an edge), thus we received a new independent set of the same cardinality and satisfying our conditions. We proceed similarly if u_1 and u_2 are not endpoints of any edge in M' and only one of v_1, v_2 is an endpoint of an edge in M' .

If none of u_1, u_2, v_1, v_2 are endpoints of edges in M' we may add uv to M' , thus contradicting the assumption M' was a largest independent set. \square

Thus we branch using the options from Lemma 3.2. There are five branches in which we choose one of the five named edges. Each edge connects two vertices of degree at least three, so in each case we obtain graph with at most $n - 6$ vertices.

Let i_1 and i_2 be the degrees of u_1 and u_2 respectively, and let j_1 and j_2 be the degrees of v_1 and v_2 in $G[V \setminus \{u_1, u_2\}]$. Assume that u_1 and u_2 share k_u neighbours, excluding u , and v_1 and v_2 share k_v neighbours, excluding v, u_1 and u_2 . The next branch is when u_1 and u_2 are taken. As we take none of the named edges, we drop the vertex u and then take one of the remaining edges leaving u_1 and one edge leaving u_2 . We have $i_1 - 1 - k_u$ ways to choose an edge leaving u_1 and $i_2 - 1 - k_u$ ways to choose one leaving u_2 . Let us now count the number of vertices removed in such branch. We remove $u_1, u_2, u, (i_1 - k_u - 1) + (i_2 - k_u - 1) + k_u$ neighbours of u_1 and u_2 , excluding u , and at least 2 neighbours of the other endpoints of the chosen edges. In total, we remove $i_1 + i_2 - k_u + 3$ vertices. Note that we do not count v to be removed, as it might be one of these 2 neighbours of the other endpoints (one of the edges taken might be, e.g., u_1v_1).

The final branch is when we drop u_1, u_2, u and v and choose one edge leaving v_1 and one edge leaving v_2 . We have $(j_1 - 1 - k_v)(j_2 - 1 - k_v)$ ways to choose these edges. From one side of the bipartite graph, we remove u_1, u_2, v and $(j_1 - 1 - k_v) + (j_2 - 1 - k_v) + k_v$ neighbours of v_1 and v_2 , excluding v, u_1 and u_2 . From the other side, we remove u and at least 2 neighbours of the other endpoints of the chosen edges. In total, we remove $j_1 + j_2 - k_v + 6$ vertices.

Therefore, in this branching rule, we obtain the following inequality for $T(n)$:

$$T(n) \geq 5T(n - 6) + (i_1 - k_u - 1)(i_2 - k_u - 1)T(n - i_1 - i_2 + k_u - 3) \\ + (j_1 - k_v - 1)(j_2 - k_v - 1)T(n - j_1 - j_2 + k_v - 6).$$

A direct check for all possible values of i_1, i_2, j_1, j_2, k_u and k_v finishes the proof that $T(n) = O(\alpha^n)$.

4 SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET

In this section we focus on solving **SMALLEST INCLUSION-MAXIMAL IRREDUNDANT SET**. We show a simple iterative-DFS algorithm that requires polynomial space and prove that it works in $O(1.999956^n)$ time. W.l.o.g. we may assume that G contains no isolated vertices, since they need to be included in any inclusion-maximal irredundant set.

Let \mathcal{F}_k be a family of irredundant sets in G of size not greater than k . Note that checking if a set is an (inclusion-maximal) irredundant set can be done in polynomial time. Moreover, the family of irredundant sets is closed under taking subsets. Therefore \mathcal{F}_k can be enumerated in $O^*(|\mathcal{F}_k|)$ time and polynomial space by simple depth-first search algorithm. The simple iterative-DFS algorithm enumerates \mathcal{F}_k for $k = 0, 1, 2, \dots, n$ until it finds a inclusion-maximal irredundant set. Now we prove that it works in $O(1.999956^n)$ time.

First, as a warm-up, let us make a $O^*((2 - \varepsilon_\Delta)^n)$ bound for graphs with maximum degree bounded by Δ , where ε_Δ depends on Δ . Construct a set $A \subset V$ greedily: repeatedly add any vertex $v \in V$ to A and remove from V all vertices distant by at most 2 from v . At each step, at most $1 + \Delta + \Delta(\Delta - 1) = 1 + \Delta^2$ vertices are removed, therefore $|A| \geq n/(1 + \Delta^2)$. The set A is an independent set; moreover, closed neighbourhoods $\{\bar{N}(v) : v \in A\}$ are disjoint. Note that any independent set S cannot contain the whole $\bar{N}(v)$ for any non-isolated vertex v , therefore:

$$|\mathcal{F}_n| \leq 2^n \prod_{v \in A} \frac{2^{|\bar{N}(v)|} - 1}{2^{|\bar{N}(v)|}} \leq 2^n \left(\frac{2^{\Delta+1} - 1}{2^{\Delta+1}} \right)^{\frac{n}{1+\Delta^2}} = (2 - \varepsilon_\Delta)^n,$$

and the time bound for the algorithm follows.

Now we show how to bypass the maximum degree assumption. Note that if an irredundant set is a dominating set, it is inclusion-maximal. Moreover an inclusion-minimal dominating set is an inclusion-maximal irredundant set. Assume that G admits a dominating set of size not greater than $149/300n$. Then the algorithm stops before or at the step $k = 149/300n$ and up to this point consumes $O^*\left(\binom{n}{149/300n}\right) = O(1.999956^n)$ time. Therefore we may consider only the case where every dominating set in G is of size greater than $149/300n$.

The following structural lemma is crucial for the analysis.




Figure 4: Notation in the proof of Lemma 4.1

Lemma 4.1. *Let $G = (V, E)$ be a graph with n vertices that contains no dominating set of cardinality smaller than $149/300n$. Then there exists a set $A \subset V$ satisfying:*

1. *A is an independent set and the neighbourhoods $\{\bar{N}(v) : v \in A\}$ are disjoint,*
2. *every vertex in A has degree at most 6,*
3. *$|A| \geq 41n/9800$.*

Proof. We construct a dominating set D greedily. Start with $D = \emptyset$. In a single step, take any vertex v that adds at least 3 new vertices to $\bar{N}(D)$, i.e. $|\bar{N}(D \cup \{v\}) \setminus \bar{N}(D)| \geq 3$, and add v to D . This algorithm stops at some point and let $A_1 = V \setminus \bar{N}(D)$, i.e. the vertices not dominated by D . For every vertex v we have $|\bar{N}(v) \cap A_1| \leq 2$, since D cannot be extended any more. In particular, every vertex in $G[A_1]$ has degree at most 1, so $G[A_1]$ is a graph of isolated vertices and isolated edges. Let A_2 be any inclusion-maximal independent set in $G[A_1]$, i.e. A_2 contains all isolated vertices of $G[A_1]$ and one endpoint of every isolated edge. A_2 is an independent set in G , too.

Let us now note that $D \cup A_2$ is a dominating set in G , since A_2 dominates A_1 . Therefore $|D| + |A_2| \geq 149/300n$. Note that, by the construction procedure of D , we have $|D| \leq \frac{1}{3}|\bar{N}(D)| = \frac{1}{3}|V \setminus A_1|$, so:

$$149/300 \leq \frac{|D| + |A_2|}{|V|} \leq \frac{1}{3} - \frac{|A_1|}{3|V|} + \frac{|A_2|}{|V|} \leq \frac{1}{3} + \frac{2}{3} \cdot \frac{|A_2|}{|V|}.$$

Therefore $|A_2| \geq 49/200n$.

Now recall that every vertex in V has at most two vertices from A_1 in its closed neighbourhood. Therefore, every vertex in V has at most two neighbours in A_2 . Let n_7 be the number of vertices in A_2 with degree at least 7. By counting edge endpoints we obtain that $7n_7 \leq 2(n - |A_2|) \leq 151/100n$ and $n_7 \leq 151/700n$. Let $A_3 \subset A_2$ be the set of vertices of degree at most 6. Then $|A_3| \geq 41/1400n$.

Now construct $A \subset A_3$ greedily. In a single step, add any $v \in A_3$ to A and remove from A_3 the vertex v and all vertices that share a neighbour with v (recall that A_3 is an independent set). Since the vertices in A_3 have degree at most 6 and every vertex in V is a neighbour of at most two vertices in A_3 , then at one step we remove at most 7 vertices from A_3 . Therefore $|A| \geq 41/9800n$. \square

The bound for our iterative-DFS algorithm is now straightforward. Note that for every non-isolated vertex v at least one point from $\bar{N}(v)$ does not belong to an irredundant set. By Lemma 4.1 we obtain $41/9800n$ disjoint sets $\{\bar{N}(v) : v \in A\}$, such that all these sets are of size at most 7 and no $\bar{N}(v)$ can be contained in an irredundant set. Therefore the total number of irredundant sets is bounded by:

$$2^n \cdot \left(\frac{2^7 - 1}{2^7} \right)^{\frac{41n}{9800}} = O(1.99994^n).$$

References

- [1] Omid Amini, Fedor V. Fomin, and Saket Saurabh. Counting subgraphs via homomorphisms. In *Proc. ICALP'09*, pages 71–82, 2009.
- [2] Andreas Björklund and Thore Husfeldt. Inclusion-exclusion algorithms for counting set partitions. In *Proc. FOCS'06*, pages 575–582, 2006.
- [3] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Fourier meets möbius: fast subset convolution. In *Proc. STOC'07*, pages 67–74, 2007.

- [4] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. The travelling salesman problem in bounded degree graphs. In *ICALP'08*, pages 198–209, 2008.
- [5] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Trimmed moebius inversion and graphs of bounded degree. In *STACS'08*, pages 85–96, 2008.
- [6] Hans Bodlaender, Daniel Lokshtanov, and Eelko Penninkx. Planar capacitated dominating set is $W[1]$ -hard. In *Proc. IWPEC'09 (to appear)*, 2009.
- [7] Nicolas Bourgeois, Bruno Escoffier, and Vangelis Th. Paschos. Efficient approximation by “low-complexity” exponential algorithms. *Cahier du LAMSADE 271, LAMSADE, Université Paris-Dauphine*, 2007.
- [8] Ljiljana Brankovic, Henning Fernau, Joachim Kneis, Dieter Kratsch, Alexander Langer, Mathieu Liedloff, Daniel Raible, and Peter Rossmanith. Private communication, 2009.
- [9] Gary Chartrand and Linda Lesniak. *Graphs and Digraphs*. CRC Press, 2004.
- [10] Marek Cygan, Lukasz Kowalik, and Mateusz Wykurz. Exponential-time approximation of weighted set cover. *Inf. Process. Lett.*, 109(16):957–961, 2009.
- [11] Marek Cygan and Marcin Pilipczuk. Faster exact bandwidth. In *Proc. WG'08*, pages 101–109, 2008.
- [12] Marek Cygan and Marcin Pilipczuk. Exact and approximate bandwidth. In *Proc. IICALP'09*, pages 304–315, 2009.
- [13] Michael Dom, Daniel Lokshtanov, Saket Saurabh, and Yngve Villanger. Capacitated domination and covering: A parameterized perspective. In *Proc. IWPEC'08*, pages 78–90, 2008.
- [14] Rodney G. Downey, Michael R. Fellows, and Venkatesh Raman. The complexity of irredundant sets parameterized by size. *Discrete Applied Mathematics*, 100(3):155–167, 2000.
- [15] Uriel Feige. Coping with the NP-hardness of the graph bandwidth problem. In *Proc. SWAT'00*, pages 10–19, 2000.
- [16] Uriel Feige and Kunal Talwar. Approximating the bandwidth of caterpillars. In *APPROX-RANDOM*, pages 62–73, 2005.
- [17] Fedor V. Fomin, Fabrizio Grandoni, and Dieter Kratsch. Measure and conquer: Domination - a case study. In *Proc. IICALP'05*, pages 191–203, 2005.
- [18] Fedor V. Fomin, Fabrizio Grandoni, and Dieter Kratsch. Measure and conquer: a simple $O(2^{0.288n})$ independent set algorithm. In *Proc. SODA'06*, pages 18–25, 2006.
- [19] Fedor V. Fomin, Kazuo Iwama, and Dieter Kratsch. Moderately exponential time algorithms, Dagstuhl seminar, 2008.

- [20] Martin Fürer, Serge Gaspers, and Shiva Prasad Kasiviswanathan. An exponential time 2-approximation algorithm for bandwidth. In *Proc. IWPEC'09, to appear*, 2009.
- [21] Marcin Mucha and Piotr Sankowski. Maximum matchings via gaussian elimination. In *Proc. FOCS'04*, pages 248–255, 2004.
- [22] Jesper Nederlof. Fast polynomial-space algorithms using möbius inversion: Improving on steiner tree and related problems. In *Proc. ICALP'09*, pages 713–725, 2009.
- [23] Venkatesh Raman and Saket Saurabh. Short cycles make $W[1]$ -hard problems hard: FPT algorithms for $W[1]$ -hard problems in graphs with no short cycles. *Algorithmica*, 52(2):203–225, 2008.
- [24] Ingo Schiermeyer. Efficiency in exponential time for domination-type problems. *Discrete Applied Mathematics*, 156(17):3291–3297, 2008.
- [25] Johan M. M. van Rooij, Jesper Nederlof, and Thomas C. van Dijk. Inclusion/exclusion meets measure and conquer: Exact algorithms for counting dominating sets. In *Proceedings of ESA'09*, pages 554–565, 2009.
- [26] Gerhard J. Woeginger. Exact algorithms for NP-hard problems: A survey. In *Combinatorial Optimization*, pages 185–208, 2001.